arXiv:1411.5389v2 [math.GR] 30 Mar 2015

UPPER BOUNDS ON THE NUMBER OF CONJUGACY CLASSES IN UNITRIANGULAR GROUPS

ANDREW SOFFER

Abstract. We provide a new upper bound on the number of conjugacy classes in the group $U_n(q)$ of unitriangular matrices over a finite field. We also compute a similar upper bound for every group in the lower central series of $U_n(q)$.

1. Introduction

For a prime power q, let $\mathbb{F}_q$ denote the field with q elements. Let $U_n(q)$ denote the group of upper-triangular matrices over $\mathbb{F}_q$ with 1s on the diagonal (called the *unitriangular group*). Let $\mathbf{k}(G)$ denote the number of conjugacy classes in a finite group G. In 1960, Higman proved the following [Hig]:

Theorem 1.1 (Higman)**.** *For every prime power q,*

$$q^{\frac{n^2}{12}(1+o(1))} \leq \mathbf{k}(U_n(q)) \leq q^{\frac{n^2}{4}(1+o(1))},$$

where $o(1)$ means a function of n, independent of q, which tends to zero as n tends to infinity.

Higman's original interest was in enumerating finite p-groups of a given order. He obtained an upper bound for the number of groups of order p^n in terms of $\mathbf{k}(U_n(p))$. While the asymptotics of the number of p-groups has since been resolved via different methods (see [BNV]), the gap between the lower and upper bounds for $\mathbf{k}(U_n(q))$ has not been closed.

In their 1992 paper, Arregi and Vera-López used their technique of canonical matrices to improve Higman's upper bound [VLA92, Theorem 5.4]. They show

$$\mathbf{k}(U_n(q)) \leq (n-1)!2^{n-1}q^{\frac{n^2+n}{6}}.$$

Note that $(n-1)!2^{n-1} = q^{O(n \log n)}$, and so these terms do not contribute significantly to the asymptotics of $\mathbf{k}(U_n(q))$. Using the theory of supercharacters (see [DI, Mar11]), Marberg obtained an upper bound with the same asymptotics [Mar08, Theorem 5.1].

We improve on these asymptotics, with the following result:

Theorem 1.2. *For every positive integer n and every prime power q, we have*

$$\mathbf{k}(U_n(q)) \leq p(n)^2 n!\ q^{\frac{7}{44}n^2+\frac{n}{2}},$$

where $p(n)$ denotes the number of integer partitions of n.

Date: April 1, 2015.

Department of Mathematics, UCLA, Los Angeles, CA 90095.
Email: asoffer@math.ucla.edu.

Our approach is to estimate the number of pairs (A, B) of commuting matrices in $U_n(q)$ by conjugating A into Jordan canonical form, and determining the possibilities for the image of B under this conjugation. There are many choices for matrices which conjugate A into Jordan form, and the image of B depends on this choice. Section 3 defines our canonical choice X_A for conjugation.

For each upper-triangular matrix A we conjugate its centralizer $C_{\mathcal{U}}(A)$ via our canonical choice X_A defined in Section 3. The resulting space $X_A C_{\mathcal{U}}(A) X_A^{-1}$ can often be described by a combinatorial object which we call a *gap array*. Section 4 introduces gap arrays and proves several structural lemmas about them. While not every space $X_A C_{\mathcal{U}}(A) X_A^{-1}$ can be encoded by a gap array, every such space is a subspace of one encoded by a gap array. Determining the sizes of these subspaces via the combinatorics of gap arrays, we obtain the same upper bound as Marberg and Arregi and Vera-López (see Corollary 5.8). However, the technique of gap arrays is amenable to further optimization. These optimizations are the content of the proof of Theorem 5.9, from which Theorem 1.2 follows (see Section 5).

In Section 6, we develop another set of tools for computing bounds on the number of conjugacy classes in groups in the lower central series for $U_n(q)$. We prove:

Theorem 1.3. *Let $U_{n,k}(q)$ denote the kth group in the lower central series for $U_n(q)$, and let*

$$\gamma_m := \frac{1}{6} - \frac{13}{24} \cdot 4^{-m} + 2^{-(m+1)} - 4^{-(m+1)} m.$$

Then for every q,

$$\mathbf{k}(U_{n,k}(q)) \leq q^{\gamma_m n^2 (1 + o_m(1))},$$

where $m = \left\lfloor \log_2 \left(\frac{n}{k} \right) \right\rfloor$, and $o_m(1)$ denotes a function which, for each fixed m, tends to zero as n tends to infinity.

The techniques we use to establish this result are entirely unrelated to gap arrays. The main idea for the proof is to bound the number of pairs of commuting matrices in $\mathcal{U}_{n,k}(q)$ by splitting each of the matrices into three parts (see Figure 4). Two of these parts are upper-triangular, and form smaller instances of the same problem. We use the fact that pairs of these smaller upper-triangular parts commute to bound the number of viable completions of the full upper-triangular matrices in terms of the ranks of the smaller parts. This bound is the content of Lemma 6.1.

We then bound the number of pairs of commuting matrices in $\mathcal{U}_{n,k}(q)$ by conditioning on the ranks of the smaller triangular parts. On the one hand, if the smaller triangular parts must have small rank, then the number of such matrices is a limiting factor. On the other hand, if the ranks of the triangular parts are large, there are stringent conditions on the viable completions, limiting the total number of possibilities.

2. Basic Definitions and Notation

2.1. **Algebra.** For a prime power q, let $\mathbb{F}_q$ denote the field with q elements, and let $\mathcal{U}_n(q)$ denote the $\mathbb{F}_q$-algebra of strictly upper-triangular matrices. Define the *unitriangular group* to be $U_n(q) = \{1 + X : X \in \mathcal{U}_n(q)\}$. The field $\mathbb{F}_q$ will almost always be clear from context. In such cases we will omit the parameter q.

For an associative algebra $\mathcal{A}$, define

$$\mathrm{Comm}(\mathcal{A}) = \{(X, Y) \in \mathcal{A} \times A : XY = YX\}.$$

For a group G, we use the same notation $\mathrm{Comm}(G)$ to denote the set of pairs of commuting elements in G. Note that $X, Y \in \mathcal{U}_n$ commute if and only if the elements $1+X, 1+Y \in U_n$ commute, so $|\mathrm{Comm}(\mathcal{U}_n)| = |\mathrm{Comm}\, U_n|$.

From Burnside's lemma, the number of conjugacy classes $\mathbf{k}(G)$ in a finite group G can be computed by

$$\mathbf{k}(G) = \frac{1}{|G|}\sum_{g\in G}|C_G(g)| = \frac{|\mathrm{Comm}(G)|}{|G|},$$

where $C_G(g)$ denotes the centralizer of $g \in G$.

For positive integers a and b, let $\mathcal{M}_{a\times b}$ denote the vector space of $(a\times b)$-matrices over $\mathbb{F}_q$. We write $C_{\mathcal{M}}(X)$ for the centralizer of X in $\mathcal{M}_{n\times n}$. That is,

$$C_{\mathcal{M}}(X) := \{A \in \mathcal{M}_{n\times n} : AX = XA\}.$$

Similarly, define $C_{\mathcal{U}}(X) := \{A \in \mathcal{U}_n : AX = XA\}$.

2.2. **Partitions.** Recall standard notation in partition theory (see, e.g. [Mac, Sta]). Let $p(n)$ denote the number of integer partitions of n. For a partition $\lambda \vdash n$, we write λ' for the conjugate partition. That is, $\lambda' = (\lambda'_1, \lambda'_2, \dots)$, where λ'_i is the number of parts in λ of size at least i. We use $\ell(\lambda)$ to denote the number of parts (or *length*) of λ. Clearly $\ell(\lambda) = \lambda'_1$.

For a partition $\lambda \vdash n$, we use the notation

$$\mathbf{n}(\lambda) := \sum_i (i-1)\lambda_i = \sum_i \binom{\lambda'_i}{2}. \tag{2.1}$$

It is common to use the character "n" for both this function and the size of a partition. We choose to use the bold roman "$\mathbf{n}$" for the function to avoid confusion.

One can also write a partition as $(1^{m_1}2^{m_2}3^{m_3}\dots)$, where m_i denotes the number of parts of size i. Clearly, m_i can be computed by $m_i = \lambda'_i - \lambda'_{i+1}$.

2.3. **Jordan canonical forms.** Let $\lambda = (\lambda_1, \dots, \lambda_\ell)$ be a partition of n. By J_λ, we denote the Jordan canonical nilpotent matrix which has blocks of size $\lambda_1, \lambda_2, \dots, \lambda_\ell$. For example,

$$J_{(3,2)} = \left(\begin{array}{ccc|cc} 0&1&0&0&0\\ 0&0&1&0&0\\ 0&0&0&0&0\\ \hline 0&0&0&0&1\\ 0&0&0&0&0 \end{array}\right).$$

Every upper-triangular matrix $A \in \mathcal{U}_n$ is GL-conjugate to some J_λ. We say that λ is the *shape* of A if A is GL-conjugate to J_λ, and write $\mathrm{sh}(A) = \lambda$. For $A \in \mathcal{U}_n$, and $1 \le k \le n$, let $A|_k$ denote the top-left $(k\times k)$-submatrix of A.

For a partiton $\lambda \vdash n$, let $F_\lambda(q)$ denote the number of matrices in $\mathcal{U}_n(q)$ of Jordan type λ. Yip proved that F_λ is a polynomial in q with integer coefficients and degree equal to $\binom{n}{2} - \mathbf{n}(\lambda)$ [Yip]. Moreover, Yip showed that the leading coefficient is f^λ, the number of standard Young tableaux of shape λ. It is clear from Yip's proof that

$$F_\lambda(q) \le f^\lambda q^{\binom{n}{2}-\mathbf{n}(\lambda)}. \tag{2.2}$$

The only fact we will use regarding standard Young tableaux is that for every $\lambda \vdash n$, we have $f^\lambda \le \sqrt{n!}$ (see, e.g. [Sag]).

2.4. **Partitions as elements of $\ell^1(\mathbb{N})$.** It will be useful to treat partitions as infinite decreasing sequences of positive integers $(\lambda_1, \lambda_2, \dots)$ where $\lambda_i = 0$ for $i > \ell(\lambda)$. In this way, we can treat λ as an element of $\ell^1(\mathbb{N})$. Moreover, as $\ell^1(\mathbb{N}) \subseteq \ell^2(\mathbb{N})$, it makes sense to talk about the inner product of two partitions. For partitions λ and μ (not necessarily of the same integer) define

$$\langle \lambda, \mu \rangle := \sum_i \lambda_i \mu_i, \text{ and}$$
$$\|\lambda\| := \sqrt{\langle \lambda, \lambda \rangle}.$$

Note that combining (2.1) with this analytic notation, the function $\mathbf{n}$ can be expressed as

$$\mathbf{n}(\lambda) = \frac{\|\lambda'\|^2}{2} - \frac{n}{2}. \tag{2.3}$$

Similarly, because the multiplicity m_i of the part i in λ can be computed as $m_i = \lambda'_i - \lambda'_{i+1}$, we have

$$(m_1, m_2, \dots) = v - Lv, \tag{2.4}$$

where L denotes the left-shift operator on $\ell^1(\mathbb{N})$.

Remark. We define $\|\cdot\|$ to be the ℓ^2-norm, rather than the ℓ^1-norm. Treating $\lambda \vdash n$ as an element of $\ell^1(\mathbb{N})$, its ℓ^1-norm is simply n, and therefore does not merit its own notation. When we need the ℓ^1-norm of a vector $v \in \ell^1(\mathbb{N})$, we will explicitly add a subscript to the norm, as in $\|v\|_1$.

3. Jordan forms and conjugation

For an upper-triangular matrix $A \in \mathcal{U}_n$, our approach to understanding the size of its centralizer $C_{\mathcal{U}}(A)$ will be to conjugate A into its Jordan form J_λ by some X_A. We then analyze $X_A C_{\mathcal{U}}(A) X_A^{-1}$, noting that it is a subspace of $C_{\mathcal{M}}(J_\lambda)$. Given $A \in \mathcal{U}_n$, there is more than one choice for a matrix X_A satisfying $X_A A X_A^{-1} = J_\lambda$. Different choices of X_A may yield different subspaces of $C_{\mathcal{M}}(J_\lambda)$. We must therefore specify X_A carefully. We do this inductively, by first assuming that $A|_{n-1} = J_\mu$ for some partition $\mu \vdash (n-1)$.

We begin by giving an overview of the process by which we put A into Jordan form. This overview, along with the example computation below, should provide enough detail for the reader to understand the conjugation process. For completeness, we also provide explicit definitions of the matrices used in the conjugation procedure.

Conjugation procedure.

(1) Use the non-zero entries from $A|_{n-1} = J_\mu$ to set as many entries as possible in column n to zero. This can be achieved with a product of upper-triangular transvections. The resulting matrix will only have non-zero entries in column n which are at the bottom of a block. That is, in cells of the form $(n, \mu_1 + \cdots + \mu_k)$. If the entire column n is zero in the result, skip ahead to step 5. In this case, the matrices by which we conjugate in the intermediate steps will all be defined as the identity.
(2) We may now assume that column n is non-zero. Conjugate by a diagonal matrix which scales the last column and last row in such a way as to set the first non-zero entry in column n to be 1.

(3) Use the first non-zero entry in column n (which now contains the value 1) to set every other value in that column to zero. This is achieved via a product of lower-triangular matrices. Each such lower triangular matrix will fix the J_μ in the top-left corner, and set a single cell in column n to be zero.
(4) Apply a permutation matrix to move column n so that its non-zero value aligns with a Jordan block, effectively increasing the size of this block by 1.
(5) At this point the matrix is in Jordan form, modulo rearranging the blocks to be in descending order. At most one block must be moved to guarantee this ordering. The block that must be moved is the one whose size was increased (if we increased the size of a block at all). Call this block the *current block*. If we did not increase the size of a block, then we created a new block of size 1. In this case, the newly created block will be called the current block. Apply a permutation matrix which moves the current block as far to the top-left as possible, while still guaranteeing that the blocks are in descending order.

Example 3.1. The following is a worked example of the conjugation procedure for a matrix $A \in \mathcal{U}_9(\mathbb{Q})$. We use $\mathbb{Q}$ as the field for simplicity, though using a finite field does not present any extra difficulty. We write $A^{[i]}$ for the matrix obtained after step i.

$$A = \begin{pmatrix} 1 & 1 & & & & & & & 3 \\ & 1 & & & & & & & 1 \\ & & & & & & & & 0 \\ \hline & & & 1 & & & & & 2 \\ & & & & & & & & 2 \\ \hline & & & & & 1 & & & 4 \\ & & & & & & & & 0 \\ \hline & & & & & & & & 1 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad A^{[1]} = \begin{pmatrix} 1 & 1 & & & & & & & 0 \\ & 1 & & & & & & & 0 \\ & & & & & & & & 0 \\ \hline & & & 1 & & & & & 0 \\ & & & & & & & & 2 \\ \hline & & & & & 1 & & & 0 \\ & & & & & & & & 0 \\ \hline & & & & & & & & 1 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad A^{[2]} = \begin{pmatrix} 1 & 1 & & & & & & & 0 \\ & 1 & & & & & & & 0 \\ & & & & & & & & 0 \\ \hline & & & 1 & & & & & 0 \\ & & & & & & & & 1 \\ \hline & & & & & 1 & & & 0 \\ & & & & & & & & 0 \\ \hline & & & & & & & & \frac{1}{2} \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$A^{[3]} = \begin{pmatrix} 1 & 1 & & & & & & & 0 \\ & 1 & & & & & & & 0 \\ & & & & & & & & 0 \\ \hline & & & 1 & & & & & 0 \\ & & & & & & & & 1 \\ \hline & & & & & 1 & & & 0 \\ & & & & & & & & 0 \\ \hline & & & & & & & & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad A^{[4]} = \begin{pmatrix} 1 & & & 1 & & & 0 & & \\ & 1 & & & & & 0 & & \\ & & & & & & 0 & & \\ \hline & & & 1 & & & 0 & & \\ & & & & & & 1 & & \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline & & & & & & 0 & 1 & \\ \hline & & & & & & 0 & & \end{pmatrix} \quad A^{[5]} = \begin{pmatrix} 1 & 0 & & & & & & & \\ & 1 & & & & & & & \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline & 0 & & 1 & & & & & \\ & 0 & & & 1 & & & & \\ & 0 & & & & & & & \\ \hline & 0 & & & & 1 & & & \\ & 0 & & & & & & & \\ \hline & 0 & & & & & & & \end{pmatrix}$$

Note that $A^{[4]} = A^{[5]}$, but the explicitly shown zeros are in different locations. This is to emphasize that the last step in our conjugation procedure may fix the Jordan form (as it does in this case). However, when we apply the same procedure to $C_{\mathcal{U}}(A)$, this last action will often be non-trivial.

Remark 3.2. It is tempting to assume that, if A is already in Jordan form, then the conjugating matrix X_A will be the identity. This is not necessarily so. While many of the steps in the conjugation procedure are trivial, the permutation matrices in steps 4 and 5 need not be the identity. For example, if $A = \mathrm{id}_n$, then X_A is the permutation matrix defined by the permutation $w : k \mapsto n + 1 - k$. It is true however that if A is already in Jordan form, then X_A will be a permutation matrix.

To be specific about this procedure, we now write down explicitly the matrices used in the conjugaction procedure. We conjugate A by a product of five matrices, one for each step in the conjugation procedure. As was shown in Example 3.1, for $i = 1, \ldots, 5$ we will use $A^{[i]}$ to denote the matrix obtained after step i.

First, conjugate by $E_A := \prod_{i=1}^{n-1} \mathcal{E}_{i+1,n}(A_{i,n})$, where $\mathcal{E}_{i,j}(\alpha)$ is the transvection having 1s on the diagonal, α in position (i,j), and zeros everywhere else. Conjugating A by E_A gives the matrix $A^{[1]} = E_A A E_A^{-1}$ satisfying

(1) $A^{[1]}|_{n-1} = J_\mu$,
(2) each row of $A^{[1]}$ has at most one non-zero entry.

Thus, the only non-zero entries in the last column of $A^{[1]}$ are in line with the bottom of a Jordan block of J_μ.

Second, let x denote the first non-zero entry in column n of $A^{[1]}$, if it exists (and $x = 1$ otherwise). Define

$$\Delta_A := \mathrm{diag}(\underbrace{1, \ldots, 1}_{n-1}, x).$$

Conjugating $A^{[1]}$ by Δ_A yields the matrix $A^{[2]} = \Delta_A A^{[1]} \Delta_A^{-1}$ which has at most one non-zero entry in each row, and which has a 1 as its first non-zero entry in column n (if column n is has any non-zero entries).

Third, define a lower-triangular matrix L_A which uses the first non-zero entry in the nth column to set all other entries in that column to zero. If column n is already zero, then simply set L_A to be the identity. Otherwise, let $\tilde{\mu}_s := \sum_{i=1}^{s} \mu_i$. These numbers are the indices of rows which are at the bottom of Jordan blocks. Every non-zero entry in column n of $A^{[2]}$ appears in such a row. Define r to be the integer such that $\tilde{\mu}_r$ is the index of the row containing the first non-zero entry in the nth column of $A^{[2]}$, and for $\alpha \in \mathbb{F}_q$ define

$$F_{j,r}(\alpha) := 1 + \alpha \sum_{k=1}^{\mu_j} e_{\tilde{\mu}_{j-1}+k,\ \tilde{\mu}_r - \mu_j + k}, \tag{3.1}$$

where $e_{i,j}$ denotes the matrix with a 1 in position (i,j) and zeros everywhere else. Left-multiplication by $F_{j,r}(\alpha)$ takes the last μ_j rows in the rth block (rows $\tilde{\mu}_r - \mu_j + 1$ through $\tilde{\mu}_r$), and adds them to the rows in the jth block (first scaling them by α). Right-multiplication by $F_{j,r}(\alpha)$ takes the μ_j columns in the jth block, and subtracts them from the last μ_j columns in the rth block (first scaling them by α). Conjugating $A^{[2]}$ by $F_{j,r}(-A^{[2]}_{n,\tilde{\mu}_j})$ will leave $A^{[2]}$ unchanged in every entry except in the entry indexed by $(n, \tilde{\mu}_j)$, which will be set to zero. We can therefore define

$$L_A := \prod_{j>r} F_{j,r}(-A^{[2]}_{n,\tilde{\mu}_j}),$$

so that $A^{[3]} = L_A A^{[2]} L_A^{-1}$ has one non-zero entry in column n, and that value is 1.

Fourth, conjugate by a permutation matrix to make column n align with the correct block. We apply $\sigma_A := (\tilde{\mu}_r + 1, \tilde{\mu}_r + 2, \ldots, n)$. If column n is zero in $A^{[3]}$ then no such permutation is necessary, and we set σ_A to be the identity. Now $A^{[4]} = \sigma_A A^{[3]} \sigma_A^{-1}$ is the direct sum of Jordan blocks, though not necessarily in descending order.

Lastly, we apply a permutation matrix τ_A which moves the block as close to the top-left as possible to put the blocks into descending order, in such a way as to preserve the relative order of all other Jordan blocks.

We are now in a position to define our choice of conjugating matrix X_A, so that $X_A A X_A^{-1}$ is in Jordan form. Define X_A recursively. For the unique $A \in \mathcal{U}_1$,

we take $X_A = (1)$ to be the identity matrix. For $n > 1$, let $B = A|_{n-1}$, and define $A' = X_B A X_B^{-1}$, so that $A'|_{n-1}$ is in Jordan form. Then define

$$Y_{A'} := \tau_{A'}\sigma_{A'}L_{A'}\Delta_{A'}E_{A'}, \text{ and}$$
$$X_A := Y_{A'}X_B.$$

Here we have implicitly identified $X_B \in \mathcal{M}_{(n-1)\times(n-1)}$ with $X_B \oplus 1 \in \mathcal{M}_{n\times n}$. By construction, we have $X_A A X_A^{-1} = J_\lambda$.

4. Gap arrays

4.1. **Preliminary results.**

Lemma 4.1. *Let $\lambda \vdash a$, $\mu \vdash b$, and let $T_{\lambda,\mu} : \mathcal{M}_{a\times b} \to \mathcal{M}_{a\times b}$ be defined by $T_{\lambda,\mu}(X) = J_\lambda X - X J_\mu$. Then*

$$\dim\ker T_{\lambda,\mu} = \sum_{i,j}\min\{\lambda_i,\mu_j\} = \langle\lambda',\mu'\rangle.$$

In particular, $\dim C_{\mathcal{M}}(J_\lambda) = \|\lambda'\|^2$.

Proof. To compute the rank, we first do so for two Jordan blocks $J_{(a)}$ and $J_{(b)}$. Note that $J_{(a)}X$ is the matrix obtained by removing the top row of X, shifting all other rows upwards by one, and adding a row of zeros at the bottom. Similarly, $XJ_{(b)}$ is the matrix obtained by removing the rightmost column of X, shifting all other columns right by one, and adding a new column of zeros on the left. If $J_{(a)}X = XJ_{(b)}$, then these conditions guarantee that

(1) all diagonals (top-left to bottom-right) are constant, and
(2) if a diagonal does not touch both the topmost row and the rightmost column, then it must be zero.

Moreover, these conditions exactly describe the solutions to $J_{(a)}X = XJ_{(b)}$. The dimension of this space is given by the quantity $\min\{a,b\}$, the number of diagonals which are not forced to be zero.

In general, if either $\lambda \vdash a$ or $\mu \vdash b$ have more than one part, then we decompose $\mathcal{M}_{a\times b}$ into smaller subspaces by $\mathcal{M}_{a\times b} \cong \bigoplus_{i,j}\mathcal{M}_{\lambda_i\times\mu_j}$, so that the action on $\mathcal{M}_{\lambda_i\times\mu_j}$ becomes $T_{(\lambda_i),(\mu_j)}$. The dimension of $\ker T_{\lambda,\mu}$ is the sum of the dimensions of the kernels of these subspaces, so

$$\dim\ker T_{\lambda,\mu} = \sum_{i,j}\min\{\lambda_i,\mu_j\}.$$

To see that $\sum_{i,j}\min\{\lambda_i,\mu_j\} = \langle\lambda',\mu'\rangle$, we will show that both sides count the number of pairs of cells in the diagrams of λ and μ (one from each diagram) which lie in the same column.

On the one hand, we may specify the column from which we chose the cells first. Then any pair of elements from these columns will suffice. Therefore, we obtain $\sum_i \lambda'_i\mu'_i = \langle\lambda',\mu'\rangle$.

On the other hand, we may first chose the row to which each of the cells belong. If the cell in the diagram for λ lies in row i, and the cell in the diagram for μ lies in row j, then there are $\min\{\lambda_i,\mu_j\}$ choices which place these cells in the same column. As these two expressions count the same quantity, we obtain

$$\sum_{i,j}\min\{\lambda_i,\mu_j\} = \langle\lambda',\mu'\rangle.$$

Lastly, taking $\lambda = \mu$, we see that $\ker T_{\lambda,\lambda}$ is the centralizer of J_λ, and therefore

$$\dim C_{\mathcal{M}}(J_\lambda) = \dim \ker T_{\lambda,\lambda} = \|\lambda'\|^2,$$

which completes the proof. □

Remark 4.2. Matrices which satisfy the first condition (constant on top-left to bottom-right diagonals) are called *Toeplitz* matrices. The matrices in $\ker T_{\lambda,\mu}$ are not only Toeplitz on each block, but are also upper-triangular. Specifically, any diagonal which does not touch the rightmost column and topmost row of a block is necessarily zero.

4.2. **Defining gap arrays.** In this section, we introduce a new combinatorial object called a *gap array*. Gap arrays are used to encode certain subspaces of $C_{\mathcal{M}}(J_\lambda)$. Many algebraic operations we can apply to these subspaces are encoded succinctly by combinatorial operations we can apply to gap arrays. This will be our primary tool for computing upper bounds on $\mathbf{k}(U_n)$.

Definition 4.3. Let $\lambda \vdash n$ be a partition of length $\ell := \ell(\lambda)$. A *gap array* of type λ is an $\ell \times \ell$ matrix $\mathbf{G} = (\mathbf{G}_{i,j})$ of integers satisfying

$$\max\{0, \lambda_i - \lambda_j\} \leq \mathbf{G}_{i,j} \leq \lambda_i. \tag{4.1}$$

We use gap arrays to define subspaces of $C_{\mathcal{M}}(J_\lambda)$. Recall that $C_{\mathcal{M}}(J_\lambda)$ has a natural block decomposition into blocks of $(\lambda_i \times \lambda_j)$-submatrices. The block with rows corresponding to the ith part of λ and columns corresponding to the jth part of λ is called the (i,j)-*block*. Specifically, it consists of all cells (x, y) in the matrix such that

$$\begin{aligned} &\lambda_1 + \cdots + \lambda_{i-1} < x \leq \lambda_1 + \cdots + \lambda_i, \text{ and} \\ &\lambda_1 + \cdots + \lambda_{j-1} < y \leq \lambda_1 + \cdots + \lambda_j. \end{aligned}$$

Definition 4.4. For a gap array $\mathbf{G}$ of type λ, we let $C(\mathbf{G})$ denote the subspace of $C_{\mathcal{M}}(J_\lambda)$ satisfying the condition that on the (i,j)-block, the lowest $\mathbf{G}_{i,j}$ diagonals touching the right boundary are all zero.

Example 4.5. Let $\lambda = (6, 2, 1, 1) \vdash 10$, and let $\mathbf{G}$ denote the gap array

$$\mathbf{G} = \begin{pmatrix} 2 & 4 & 6 & 5 \\ 1 & 0 & 1 & 2 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

Then the corresponding subspace $C(\mathbf{G})$ of $C_{\mathcal{M}}(J_\lambda)$ consists of all 10×10 matrices of the form shown below (zeros omitted). In this example, $C(\mathbf{G})$ is a 16-dimensional subspace of $C_{\mathcal{M}}(J_\lambda)$.

$$\left(\begin{array}{ccccc|cc|c|c} & a_3 & a_4 & a_5 & a_6 & b_5 & b_6 & & c_6 \\ & & a_3 & a_4 & a_5 & & b_5 & & \\ & & & a_3 & a_4 & & & & \\ & & & & a_3 & & & & \\ \hline & & & & d_2 & e_1 & e_2 & f_2 & \\ & & & & & & e_1 & & \\ \hline & & & & g_1 & & & & h_1 \\ \hline & & & & & & i_1 & j_1 & k_1 \end{array}\right)$$

Remark 4.6. The name "gap array" is chosen, because the recorded numbers measure the gap between the bottom of a block and the triangle it contains. Recall from the proof of Lemma 4.1 that a $C_{\mathcal{M}}(J_\lambda)$ is characterized by, being Toeplitz on each block, and having any diagonal which does not touch both the topmost row and rightmost column of a block set to zero. Thus, $C_{\mathcal{M}}(J_\lambda)$ can be expressed as $C(\mathbf{G})$, where $\mathbf{G}_{i,j} = \max\{0, \lambda_i - \lambda_j\}$. At the other extreme, the zero subspace of $\mathcal{M}_{n\times n}$ is $C(\mathbf{G})$, where $\mathbf{G}_{i,j} = \lambda_i$. The inequalities in (4.1) which defines gap arrays encode these observations.

For gap arrays $\mathbf{G} = (\mathbf{G}_{i,j})$ and $\mathbf{H} = (\mathbf{H}_{i,j})$ of the same type, we say that $\mathbf{G} \leq \mathbf{H}$ if $\mathbf{G}_{i,j} \leq \mathbf{H}_{i,j}$ for all i and j. It follows that $\mathbf{G} \leq \mathbf{H}$ if and only if $C(\mathbf{H}) \subseteq C(\mathbf{G})$. We also write $|\mathbf{G}| = \sum_{i,j} \mathbf{G}_{i,j}$.

Proposition 4.7. *Let $\mathbf{G}$ be a gap array of type $\lambda \vdash n$, and let $\ell = \ell(\lambda)$. Then*

$$\dim C(\mathbf{G}) = n\ell - |\mathbf{G}|\,.$$

Proof. Observe that $\mathbf{G}_{i,j}$ counts the number of diagonals touching the rightmost column of the (i,j)-block which are zero in every element of $C(\mathbf{G})$. Thus, $|\mathbf{G}|$ counts the number of diagonals in all blocks which are necessarily zero. The quantity $\dim C(\mathbf{G})$ counts the number of diagonals in all blocks which may be non-zero. This constitutes all diagonals, so $|\mathbf{G}| + \dim C(\mathbf{G})$ is a constant depending only on the type λ of $\mathbf{G}$. Taking $\dim C(\mathbf{G}) = 0$, by setting $\mathbf{G}_{i,j} = \lambda_i$, we see that $|\mathbf{G}| = \sum_{i,j} \lambda_i = n\ell$. Hence, for every gap arary $\mathbf{G}$ of type λ,

$$\dim C(\mathbf{G}) = n\ell - |\mathbf{G}|\,. \tag{4.2}$$

□

4.3. **Combinatorial operations on gap arrays.** Recall from Section 3 that we can put a matrix A into Jordan form by iteratively conjugating larger and larger submatrices into Jordan form. As we do so, we conjugate $C_{\mathcal{U}}(A)$ by the same process. The purpose of this section is to translate what happens to $C_{\mathcal{U}}(A)$ in this iterative procedure into the combinatorics of gap arrays.

For a partition λ, and for $r \in \{1, \ldots, \ell(\lambda)\}$, let $\phi_r(\lambda)$ denote the partiton obtained by adding a block to the diagram for λ into row r. If $\lambda_r = \lambda_{r-1}$, then the result will be a composition and no longer a partition, so we reorder the rows to obtain a partition. For example, if $\lambda = (3^1 2^2 1^4)$, then $\phi_6(\lambda) = (3^1 2^3 1^3)$. If $r = \ell(\lambda) + 1$, then we define $\phi_r(\lambda)$ to be the partition obtained by adding a part of size 1.

For a gap array $\mathbf{G}$ of type $\lambda = (\lambda_1, \ldots, \lambda_\ell)$, we define a new gap array $\psi_r(\mathbf{G})$ of type $\phi_r(\lambda)$. Define $\psi_r(\mathbf{G})$ as follows:

(1) If $r = \ell + 1$, add a new row of all zeros at the bottom of $\mathbf{G}$, and a new column at the right of $\mathbf{G}$ with the values $(\lambda_1, \ldots, \lambda_\ell, 0)$.
(2) Subtract 1 from each non-zero entry in column r of $\mathbf{G}$.
(3) Add 1 to each entry in row r.
(4) Permute the rows and columns so that corresponding block sizes are in decreasing order. Specifically, move row and column r to be the very first row and column of their particular block size. Keep all other rows/columns in the same relative order.

Proposition 4.8. *Let $\mathbf{G}$ be a gap array of type λ, and let $r \in \{1, \ldots, \ell(\lambda) + 1\}$. Then $\psi_r(\mathbf{G})$ is a gap array of type $\phi_r(\lambda)$.*

Proof. First note that to obtain $\psi_r(\mathbf{G})$, the rows and columns of $\mathbf{G}$ are rearranged according to the same permutation as the entries in λ when computing $\phi_r(\lambda)$. We can therefore ignore the permutations and check that that the entries in $\psi_r(\mathbf{G})$ satisfy the inequalities given in (4.1).

Note that $\mathbf{G}_{i,j}$ will be increased (by one) if and only if $i = r$. Because λ_r is also increased by one, the upper bound given in (4.1) will still be satisfied. From the definition of ψ_r, we see that no entry in $\psi_r(\mathbf{G})$ can be negative. Hence, the only way an inequality from (4.1) can fail to be satisfied is if $\mathbf{G}_{i,j}$ is decreased (by one). However, in this situation, it must be that $j = r$, and so the lower bound of $\lambda_i - \lambda_j$ is also decreased by one. Thus, the inequalities from (4.1) are always preserved, and so $\psi_r(\mathbf{G})$ is a gap array of type $\phi_r(\mathbf{G})$. □

Example 4.9. Starting with $\lambda = (3, 2, 1) \vdash 6$, we apply ϕ_2, ϕ_4, and ϕ_4 in that order. We obtain

$$\begin{aligned} \phi_2(\lambda) &= (3, 3, 1) \\ \phi_4(\phi_2(\lambda)) &= (3, 3, 1, 1) \\ \phi_4(\phi_4(\phi_2(\lambda))) &= (3, 3, 2, 1). \end{aligned}$$

Correspondingly, for a gap array, let

$$\mathbf{G} = \begin{pmatrix} 1 & 1 & 2 \\ 1 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}.$$

Applying ψ_2, ψ_4, and ψ_4 in that order to $\mathbf{G}$, we obtain

$$\psi_2(\mathbf{G}) = \begin{pmatrix} 1 & 2 & 2 \\ 0 & 1 & 2 \\ 0 & 0 & 1 \end{pmatrix}, \quad \psi_4(\psi_2(\mathbf{G})) = \begin{pmatrix} 1 & 2 & 2 & 2 \\ 0 & 1 & 2 & 2 \\ 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \text{ and}$$

$$\psi_4(\psi_4(\psi_2(\mathbf{G}))) = \begin{pmatrix} 1 & 2 & 1 & 2 \\ 0 & 1 & 1 & 2 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 \end{pmatrix}.$$

Definition 4.10. Let $\lambda \vdash n$, and let $r \in \{1, \ldots, \ell(\lambda)\}$. Then a gap array $\mathbf{G}$ of type λ is called *r-valid* if the following two conditions hold

(1) Row r is element-by-element weakly larger than every row below it. That is, for every $j > r$, and every k,

$$\mathbf{G}_{j,k} \leq \mathbf{G}_{r,k}.$$

(2) Column r is element-by-element weakly smaller than every column to the right. That is, for every $j > r$ and every k,

$$\mathbf{G}_{k,r} \leq \mathbf{G}_{k,j}.$$

Validity is a technical condition which will be needed in Lemma 4.11. The most important property to recognize is that a gap array is r-valid for every r if and only if its entries are increasing from left to right and from bottom to top.

We need one more definition to state Lemma 4.11. For a subspace $V \subseteq \mathcal{M}_{(n-1)\times(n-1)}$, define

$$\overline{V} := \{X \in \mathcal{M}_{n\times n} : X|_{n-1} \in V, X_{n,i} = 0 \text{ for all } i \leq n-1\}.$$

Graphically, we are considering the subspace shown in Figure 1. We are now ready to state and prove our main lemma regarding gap arrays.

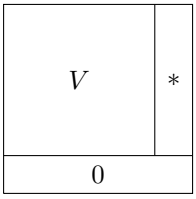

FIGURE 1. A graphical representation of $\overline{V}$, where $V \subseteq \mathcal{M}_{(n-1)\times(n-1)}$.

Lemma 4.11. *Let $A \in \mathcal{U}_n$ and let $\mu \vdash (n-1)$ such that $A|_{n-1} = J_\mu$. Let $\mathbf{G}$ be an r-valid gap array of type μ. Then*

$$Y_A \left(\overline{C(\mathbf{G})} \cap C_{\mathcal{M}}(A)\right) Y_A^{-1} = C(\psi_r(\mathbf{G})).$$

Proof. We first consider $Y_A\overline{C(\mathbf{G})}Y_A^{-1}$. Note that for any subspace $V \subseteq \mathcal{M}_{(n-1)\times(n-1)}$, the matrix E_A normalizes $\overline{V}$, as its action only changes the last column, and fixes the bottom-right entry. Similarly, Δ_A normalizes $\overline{V}$. Hence,

$$Y_A\overline{C(\mathbf{G})}Y_A^{-1} = (\tau_A\sigma_A L_A)\overline{C(\mathbf{G})}(\tau_A\sigma_A L_A)^{-1}.$$

Next, we claim that L_A normalizes $\overline{C(\mathbf{G})}$. Let $A_{[i,j]}$ denote the $(\mu_i \times \mu_j)$-matrix obtained by restricting A to the (i,j)-block. Recall that L_A is defined in (3.1) to be the product of the lower-triangular matrices of the form $F_{j,r}(\alpha)$. Let $B \in \overline{C(\mathbf{G})}$. Then by construction, $F_{j,r}(\alpha) \cdot B$ agrees with B on all blocks except for those of the form $B_{[j,k]}$ for some k. On such blocks, the action of left-multiplication by $F_{j,r}(\alpha)$ is to add the bottom λ_j rows of $\alpha \cdot B_{[r,k]}$ to $B_{[j,k]}$. Because $\mathbf{G}$ is r-valid, and $j > r$, we know that $\mathbf{G}_{j,k} \leq \mathbf{G}_{r,k}$ for all k. Thus, every diagonal of non-zero entries in $\alpha \cdot B_{[r,k]}$ gets added to a diagonal in $B_{[j,k]}$ which is allowed to be non-zero. As a result, left-multiplication by $F_{j,r}(\alpha)$ stabilizes $\overline{C(\mathbf{G})}$.

This is shown graphically by the left diagram in Figure 2. The light-gray strips represent the rows and columns of the relevant blocks. The vertical strip represents the $(*,k)$-blocks, and the horizontal strips represent the $(r,*)$- and $(j,*)$-blocks. The dark-gray triangles represent those diagonals in each block which may be non-zero. The black triangle shows the cells in the (r,k)-block which get carried to the (j,k)-block via the left-multiplication by $F_{j,r}(\alpha)$.

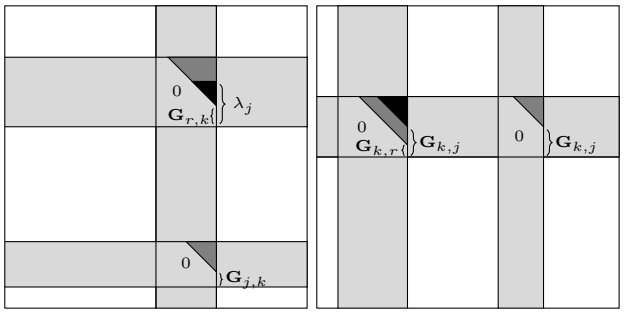

FIGURE 2. A graphical representation of how r-validity implies that conjugation by $F_{j,r}(\alpha)$ stabilizes $\overline{C(\mathbf{G})}$.

Similarly, right-multiplication by $F_{j,r}(\alpha)^{-1}$ maps $\overline{C(\mathbf{G})}$ to itself as a consequence of the column-condition of r-validity. For each k, the action of right-multiplication by $F_{j,r}(\alpha)^{-1}$ subtracts $\alpha \cdot B_{[k,j]}$ from $B_{[k,r]}$, where the right edges of each block are aligned, and $\alpha \cdot B_{[k,j]}$ is extended to the left with zeros if necessary. Because $\mathbf{G}$ is r-valid, and $j > r$, we know that $\mathbf{G}_{k,r} \leq \mathbf{G}_{k,j}$ for all k. Thus, every diagonal of non-zero entries in $B_{[k,j]}$ gets subtracted from a diagonal in $B_{[k,r]}$ which is allowed to be non-zero.

This is shown graphically by the diagram on the right in Figure 2. The dark-gray triangles represent those diagonals in each block which may be non-zero. The black triangle shows where the right-multiplication by $F_{j,r}(\alpha)^{-1}$ carries the triangle from the (k,j)-block. Because $\mathbf{G}_{k,j} \geq \mathbf{G}_{k,r}$, the black triangle lies inside the dark-gray triangle in the (k,r)-block. Hence, right-multiplication by $F_{j,r}(\alpha)^{-1}$ stabilizes $\overline{C(\mathbf{G})}$. In this way, we see that r-validity is a combinatorial description of the fact that $F_{j,r}(\alpha)$ normalizes $\overline{C(\mathbf{G})}$. It follows that $Y_A\overline{C(\mathbf{G})}Y_A^{-1} = (\tau_A\sigma_A)\overline{C(\mathbf{G})}(\tau_A\sigma_A)^{-1}$.

Next, recognize that σ_A moves the last row and column into the rth block, effectively increasing the size of the rth block by 1. Then τ_A acts by rearranging these blocks, guaranteeing that their sizes are in descending order. These permutations have the same action on the block sizes as the function ϕ_r does to the parts of μ. Thus, the block sizes are now described by the partition $\phi_r(\mu)$. Hence,

$$Y_A\left(\overline{C(\mathbf{G})} \cap C_{\mathcal{M}}(A)\right)Y_A^{-1} = (\tau_A\sigma_A)\overline{C(\mathbf{G})}(\tau_A\sigma_A)^{-1} \cap C_{\mathcal{M}}(J_{\phi_r(\mu)}). \tag{4.3}$$

The left-hand side of (4.3) is a subspace of $C_{\mathcal{M}}(J_{\phi_r(\mu)})$. Therefore every block must be Toeplitz (constant on diagonals), and any diagonal not touching both the topmost row and rightmost column of a block must be zero. We now consider how the action of $\tau_A\sigma_A$ affects each block. Because τ_A only permutes the blocks, we focus our attention on the action of σ_A. If $i \neq r$ and $j \neq r$, then σ_A has no affect on an (i,j)-block. This coincides with the fact that ψ_r does not affect cells (i,j) in a gap array when $i \neq r$ and $j \neq r$.

Considering an (r,j)-block for $j \neq r$, we see that σ_A moves the last row to the bottom of such a block. As this row is necessarily zero, it increases the gap on an (r,j)-block by 1. This is precisely what the map ψ_r encodes by adding 1 to each entry in row r of the gap array.

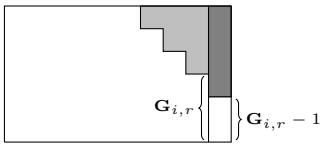

FIGURE 3. A graphical representation of an (i,r)-block is affected by the action of σ_A (for $i \neq r$).

Considering an (i,r)-block, for $i \neq r$, we see that σ_A moves the last column to the right edge of this block. Because the result is guaranteed to be Toeplitz, the size of the gap can only decrease by 1, as shown in Figure 3. If the gap size on an (i,r)-block was already zero in $C(\mathbf{G})$, then such matrices cannot possibly lie in $Y_A\left(\overline{C(\mathbf{G})} \cap C_{\mathcal{M}}(A)\right) Y_A^{-1}$, because the action of σ_A would force the (i,r)-block to no longer satisfy the upper-triangular conditions necessary to lie in $C_{\mathcal{M}}(J_{\phi_r(\mu)})$. Thus, the resulting block will have its gap size decreased by 1 unless it is already zero, in which case, the gap size remains zero. This is encoded by the action of ψ_r on the rth column of $\mathbf{G}$.

Regarding the (r,r)-block, both its width and height are increased by 1. As the resulting space is guaranteed to be Toeplitz, the gap size does not change. This is consistent with $\psi_r(\mathbf{G})$. Thus, it follows that $Y_A\left(\overline{C(\mathbf{G})} \cap C_{\mathcal{M}}(A)\right) = C(\psi_r(\mathbf{G}))$ as desired. □

5. Proof of Theorem 1.2

5.1. **Overview of proof methods for Theorem 1.2.** In this section, we prove Theorem 5.9, which trivially implies Theorem 1.2. Our proof proceeds as follows. We first construct, for each $\lambda \vdash n$, a gap array $\mathbf{G}^\lambda$ that encodes a particular subspace of $C_{\mathcal{M}}(J_\lambda)$. The space $C(\mathbf{G}^\lambda)$ contains every space of the form $X_A C_{\mathcal{U}}(A) X_A^{-1}$, where $\mathrm{sh}(A) = \lambda$. This is statement is made precise in Theorem 5.5. The gap array $\mathbf{G}^\lambda$ is particularly easy to analyze combinatorially. We compute $\left|\mathbf{G}^\lambda\right|$ in Proposition 5.4. We then show that the worst possible choice λ gives an exponent of $\frac{1}{6}n^2$. The content of the proof of Theorem 5.9 is to combine this technique with a simpler bound in order to improve this exponent.

Recall the notation $\mathrm{Comm}(\mathcal{U}_n) = \{(A,B) \in \mathcal{U}_n \times \mathcal{U}_n : AB = BA\}$. We introduce two families of subspaces of $\mathrm{Comm}(\mathcal{U}_n)$ which we parameterize with partitions. For partitions $\lambda, \mu \vdash n$, define

$$\mathrm{Comm}(\lambda) := \{(A,B) \in \mathcal{U}_n \times \mathcal{U}_n : AB = BA,\ \mathrm{sh}(A) = \lambda\},\ \text{and}$$
$$\mathrm{Comm}(\lambda,\mu) := \{(A,B) \in \mathcal{U}_n \times \mathcal{U}_n : AB = BA,\ \mathrm{sh}(A) = \lambda,\ \mathrm{sh}(B) = \mu\}.$$

We improve on the $\frac{1}{6}n^2$ exponent by first proving a technical lemma (Lemma 5.7) regarding the bounds on each $\mathrm{Comm}(\lambda)$. This lemma allows us to determine when gap arrays will be useful, and when we should bound $\mathbf{k}(U_n)$ via different, but simpler, techniques.

5.2. **Bounds on gap arrays.** For any partition $\lambda \vdash n$, there is a specific gap array whose corresponding subspace of $C_{\mathcal{M}}(J_\lambda)$ contains all possible subspaces coming from upper-triangular matrices. Moreover, we can construct this gap array explicitly

Definition 5.1. For $\lambda \vdash n$, we define the gap array $\mathbf{G}^\lambda$ by

$$\mathbf{G}^\lambda_{i,j} = \begin{cases} \lambda_i - \lambda_j & \text{if } \lambda_i > \lambda_j \\ 1 & \text{if } \lambda_i = \lambda_j \text{ and } i \leq j \\ 0 & \text{otherwise.} \end{cases}$$

Example 5.2. For the partition $\lambda = (6, 3, 1, 1, 1)$, we have

$$\mathbf{G}^\lambda = \begin{pmatrix} 1 & 3 & 5 & 5 & 5 \\ 0 & 1 & 2 & 2 & 2 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Proposition 5.3. *For every $\lambda \vdash n$, the gap array $\mathbf{G}^\lambda$ is r-valid for every r, and*

$$\mathbf{G}^{\phi_r(\lambda)} \leq \psi_r(\mathbf{G}^\lambda).$$

Proof. A gap array is r-valid for every r if and only if the entries are weakly increasing from left to right and from bottom to top. From the definition of $\mathbf{G}^\lambda$, it is clear that $\mathbf{G}^\lambda$ satisfies this property. It therefore suffices to prove that $\mathbf{G}^{\phi_r(\lambda)} \leq \psi_r(\mathbf{G}^\lambda)$.

First suppose that $r = \ell + 1$, where $\ell = \ell(\lambda)$. Recalling the definition of ψ_r from Section 4.3, we see that that $\psi_r(\mathbf{G}^\lambda)$ adds a new column containing the value $\lambda_i - 1$ in entry $(i, \ell + 1)$, and a new row of all 1s in row $\ell + 1$. This row and column are then permuted to be the first row and column corresponding to a part of $\phi_r(\lambda)$ of size 1. We may obtain $\mathbf{G}^{\phi_r(\lambda)}$ can be obtained from $\mathbf{G}^\lambda$ by adding a new column containing the number $\lambda_i - 1$ in the cell associated to the ith part of λ, and a new row containing 1s and 0s. The new columns and rows in these two gap arrays align by construction. If the new values in these rows and columns do not agree, it is because a cell has the value 0 in $\mathbf{G}^{\phi_r(\lambda)}$, and the value 1 in $\psi_r(\mathbf{G}^\lambda)$. Thus, we obtain $\mathbf{G}^{\phi_r(\lambda)} \leq \psi_r(\mathbf{G}^\lambda)$.

Now suppose $1 \leq r \leq \ell(\lambda)$. The map ψ_r acts by subtracting 1 from each non-zero entry in column r, adding 1 to each entry in row r, and then permuting the rows and columns so that this row and column are the first corresponding to a part of size $\lambda_r + 1$. Let w denote this permutation. Let s denote the new index of row/column r after this permutation, and let $\mathbf{H} = \psi_r(\mathbf{G}^\lambda)$. Then from the definition of the map ψ_r in Section 4.3, we see that

(1) $\mathbf{H}_{w(r),w(r)} = 1$,
(2) $\mathbf{H}_{w(r),w(j)} = \mathbf{G}^\lambda_{r,j} + 1$ for $j \neq r$,
(3) $\mathbf{H}_{w(i),w(r)} = \max\{0, \mathbf{G}^\lambda_{i,r} - 1\}$ for $i \neq r$,
(4) $\mathbf{H}_{w(i),w(j)} = \mathbf{G}^\lambda_{i,j}$ for all $i, j \neq r$.

Now consider $\mathbf{G}^{\phi_r(\lambda)}$. Recall that ϕ_r acts on λ by increasing the rth part, and then permuting the parts so that the resulting composition becomes a permutation again. Specifically, this can be done by applying the same permutation w as above. In other words, let $\mu := \phi_r(\lambda)$. Then

$$\mu_{w(i)} = \begin{cases} \lambda_r + 1 & \text{if } i = r \\ \lambda_i & \text{if } i \neq r. \end{cases}$$

We now show that for all i and j, we have $\mathbf{G}^\mu_{w(i),w(j)} \leq \mathbf{H}_{w(i),w(j)}$. We break this into four cases to coincide with the four items defining $\mathbf{H}_{w(i),w(j)}$ above.

Case 1: $i = j = r$**.**
It is immediate from the definition of $\mathbf{G}^\mu$ that every entry on the main diagonal is equal to 1. Thus,

$$\mathbf{G}^\mu_{w(r),w(r)} = 1 = \mathbf{H}_{w(r),w(r)}.$$

Case 2: $i = r$**,** $j \neq r$**.**
If $\mu_{w(r)} > \mu_{w(j)}$, then $\mathbf{G}^\mu_{w(r),w(j)} = \mu_{w(r)} - \mu_{w(j)}$. We therefore obtain

$$\begin{aligned}\mathbf{G}^\mu_{w(r),w(j)} &= \mu_{w(r)} - \mu_{w(j)} \\ &= \lambda_r + 1 - \lambda_j = \mathbf{G}^\lambda_{r,j} + 1 = \mathbf{H}^\lambda_{w(r),w(j)}.\end{aligned}$$

Otherwise, we see that $\mathbf{G}^\mu_{w(r),w(j)} \leq 1$, and $\mathbf{H}_{w(r),w(j)} = \mathbf{G}^\lambda_{r,j} + 1 \geq 1$. Thus, we have $\mathbf{G}^\mu_{w(r),w(j)} \leq \mathbf{H}_{w(r),w(j)}$ as desired.

Case 3: $i \neq r$**,** $j = r$**.**
First suppose $\mu_{w(i)} > \mu_{w(r)}$, so that $\mathbf{G}^\mu_{w(i),w(r)} = \mu_{w(i)} - \mu_{w(r)}$. Thus, we see that $\lambda_i - (\lambda_r + 1) > 0$, so $\mathbf{H}_{w(i),w(r)} = \mathbf{G}^\lambda_{i,r} - 1$. We therefore obtain

$$\mathbf{G}^\mu_{w(i),w(r)} = \mu_{w(i)} - \mu_{w(r)} = \lambda_i - (\lambda_r + 1) = \mathbf{H}_{w(i),w(r)}.$$

Next, suppose that $\mu_{w(i)} = \mu_{w(r)}$, and $w(i) \leq w(r)$. Recall that w is defined to have the property that $w(r)$ is the first part in μ of size $\mu_{w(r)}$. It follows that if $\mu_{w(i)} = \mu_{w(r)}$ and $w(i) \leq w(r)$, then in fact $w(i) = w(r)$. This is not possible, as $i \neq r$ by hypothesis.

In the remaining cases, $\mathbf{G}^\mu_{w(i),w(r)} = 0$. We know from the definition of gap arrays that $\mathbf{H}_{w(i),w(r)} \geq 0$ as desired.

Case 4: $i \neq r$**,** $j \neq r$**.**
Recalling that for $i \neq r$, we have $\mu_{w(i)} = \lambda_i$, we see that

$$\mathbf{G}^\mu_{w(i),w(j)} = \mathbf{G}^\lambda_{i,j} = \mathbf{H}_{w(i),w(j)}.$$

Thus, for every i and j, we have $\mathbf{G}^\mu_{w(i),w(j)} \leq \mathbf{H}_{w(i),w(j)}$, so we obtain the relationship $\mathbf{G}^{\phi_r(\lambda)} \leq \psi_r(\mathbf{G}^\lambda)$ as desired. □

Proposition 5.4. *For any $\lambda \vdash n$ with $\ell = \ell(\lambda)$, we have*

$$\left|\mathbf{G}^\lambda\right| = n\ell - n - 2\,\mathbf{n}(\lambda) + \sum_i \frac{m_i^2}{2} + \frac{\ell}{2}.$$

Proof. Recall that $\mathbf{G}^\lambda$ is an upper-triangular array, and that for $i \leq j$ and $\lambda_i = \lambda_j$, we have $\mathbf{G}^\lambda_{i,j} = 1$. Such cells in $\mathbf{G}^\lambda$ contribute $\sum_i \binom{m_i+1}{2}$. For $\lambda_i > \lambda_j$, we have $\mathbf{G}^\lambda_{i,j} = \lambda_i - \lambda_j$. Altogether, we have

$$\left|\mathbf{G}^\lambda\right| = \sum_{i>j} \lambda_i - \sum_{i>j} \lambda_j + \sum_i \binom{m_i+1}{2}.$$

To compute $\sum_{i>j} \lambda_i$, we instead sum over all possible pairs (i, j), and subtract those for which $i \leq j$. Summing the term λ_i over all possible pairs (i, j) we obtain $n\ell$. Those terms where $i = j$ contribute $\sum_i \lambda_i = n$. Thus, we have

$$\left|\mathbf{G}^\lambda\right| = n\ell - n - 2\sum_{i>j} \lambda_j + \sum_i \frac{m_i^2}{2} + \sum_i \frac{m_i}{2}.$$

From (2.1), we see that $\sum_{i>j} \lambda_j = \sum_j (j-1)\lambda_j = \mathbf{n}(\lambda)$. Furthermore, $\sum_i m_i = \ell$, hence

$$\left|\mathbf{G}^\lambda\right| = n\ell - n - 2\,\mathbf{n}(\lambda) + \sum_i \frac{m_i^2}{2} + \frac{\ell}{2}.$$

□

Theorem 5.5. *For any $A \in \mathcal{U}_n$ with Jordan form J_λ, we have*

$$X_A C_{\mathcal{U}}(A) X_A^{-1} \subseteq C(\mathbf{G}^\lambda).$$

Proof. We proceed by induction on n. In the base case, when $n = 1$, we see that A must be the 1×1 zero-matrix, and X_A is the 1×1 identity matrix. Thus, $X_A C_{\mathcal{U}}(A) X_A^{-1}$ consists of only the zero matrix, which is equal to $C(\mathbf{G}^{(1)})$.

For $n > 1$, let $A \in \mathcal{U}_n$ be GL-conjugate to J_λ. Let $B := A|_{n-1}$ be GL-conjugate to J_μ. The key to our inductive step is the equality

$$C_{\mathcal{U}}(A) = \overline{C_{\mathcal{U}}(B)} \cap C_{\mathcal{M}}(A). \tag{5.1}$$

To see this equality, first note that for two upper-triangular matrices X and Y to commute, it must also be that $X|_k$ and $Y|_k$ commute for every k. Hence, we have $C_{\mathcal{U}}(A) \subseteq \overline{C_{\mathcal{U}}(B)}$. It is immediate that $C_{\mathcal{U}}(A) \subseteq C_{\mathcal{M}}(A)$, proving that the left-hand side is contained in the right-hand side. In the other direction, observe that $\overline{C_{\mathcal{U}}(B)}$ consists only of upper-triangular matrices, and that $C_{\mathcal{M}}(A)$ consists only of matrices that commute with A.

We implicitly embed X_B into $\mathcal{M}_{n\times n}$ by identifying X_B with $X_B \oplus 1$. With this identification, we have

$$X_B \overline{V} X_B^{-1} = \overline{X_B V X_B^{-1}}, \text{ for any subspace } V \subseteq \mathcal{M}_{(n-1)\times(n-1)}.$$

Now consider conjugating both sides of (5.1) by X_B. We obtain

$$X_B C_{\mathcal{U}}(A) X_B^{-1} \;\subseteq\; \overline{X_B C_{\mathcal{U}}(B) X_B^{-1}} \;\cap\; (X_B C_{\mathcal{M}}(A) X_B^{-1}).$$

By inductive hypothesis, we may assume that $X_B C_{\mathcal{U}}(B) X_B^{-1} \subseteq C(\mathbf{G}^\mu)$. Thus,

$$X_B C_{\mathcal{U}}(A) X_B^{-1} \subseteq \overline{C(\mathbf{G}^\mu)} \cap C_{\mathcal{M}}(A'),$$

where $A' = X_B A X_B^{-1}$. Now A' satisfies the hypotheses of Lemma 4.11, so we may further conjugate by $Y_{A'}$ to obtain

$$Y_{A'} X_B C_{\mathcal{U}}(A) X_B^{-1} Y_{A'}^{-1} \subseteq Y_{A'} \left(\overline{C(\mathbf{G}^\mu)} \cap C_{\mathcal{M}}(A')\right) Y_{A'}^{-1} = C(\psi_r(\mathbf{G}^\mu)).$$

From Proposition 5.3, we know that $\mathbf{G}^\lambda \leq \psi_r(\mathbf{G}^\mu)$, from which it follows immediately that $C(\psi_r(\mathbf{G}^\mu)) \subseteq C(\mathbf{G}^\lambda)$ as desired. □

5.3. **Bounds on $\mathbf{k}(U_n)$.** We will bound $\mathbf{k}(U_n)$ by using a map $h : \ell^2(\mathbb{N}) \to \mathbb{R}$. Specifically, define $h(v) := \|v\|^2 - \|v - Lv\|^2$. The following theorem highlights how we use this function.

Theorem 5.6. *For every $\lambda \vdash n$, we have*

$$|\mathrm{Comm}(\lambda)| \leq \sqrt{n!}\, q^{\frac{1}{2}(n^2 + h(\lambda'))}.$$

Proof. Observe that

$$|\mathrm{Comm}(\lambda)| = \sum_{A:\mathrm{sh}(A)=\lambda} |C_{\mathcal{U}}(A)|\,.$$

From Theorem 5.5 and Proposition 4.7, we know that for every A,

$$|C_{\mathcal{U}}(A)| \leq \left|C(\mathbf{G}^\lambda)\right| = q^{n\ell - \left|\mathbf{G}^\lambda\right|}.$$

Combining this fact with (2.2), we see that

$$|\mathrm{Comm}(\lambda)| \leq F_\lambda\ q^{n\ell - \left|\mathbf{G}^\lambda\right|} \leq f^\lambda q^{\binom{n}{2} - \mathbf{n}(\lambda) + n\ell - \left|\mathbf{G}^\lambda\right|},$$

It is a basic result from the representation theory of the symmetric group that $f^\lambda \leq \sqrt{n!}$. It therefore suffices to show that $\binom{n}{2} - \mathbf{n}(\lambda) + n\ell - \left|\mathbf{G}^\lambda\right| \leq \frac{1}{2}(n^2 + h(\lambda'))$, or equivalently, that

$$n\ell - \tfrac{n}{2} - \mathbf{n}(\lambda) - \left|\mathbf{G}^\lambda\right| \leq \tfrac{1}{2}h(\lambda'). \tag{5.2}$$

Combining (2.3), (2.4), and Proposition 5.4 with the left-hand side of (5.2), we obtain

$$\begin{aligned} n\ell - \frac{n}{2} - \mathbf{n}(\lambda) - \left|\mathbf{G}^\lambda\right| &= \frac{n}{2} - \frac{\ell}{2} + \mathbf{n}(\lambda) - \sum_i \frac{m_i^2}{2} \\ &= -\frac{\ell}{2} + \frac{1}{2}\left(\|\lambda'\|^2 - \|\lambda' - L\lambda'\|^2\right) \\ &\leq \frac{1}{2}h(\lambda'), \end{aligned}$$

which completes the proof. □

Lemma 5.7. *Let $D \subseteq \ell^1(\mathbb{N})$ denote the vector space of all $v \in \ell^1(\mathbb{N})$ satisfying the condition[1] that $v_1 \geq v_2 \geq \cdots \geq 0$. Then the following hold:*

1. *For every $c \in \mathbb{R}$, $h(cv) = c^2 h(v)$.*
2. *$\frac{dh}{dv_1} \geq \frac{dh}{dv_k}$ for all positive integers $k \geq 3$.*
3. *For all $v \in D$, we have $h(v) \leq 2\|v\|_1 v_2 - 3v_2^2$.*
4. *For all $v \in D$ satisfying $v_1 \geq \frac{1}{2}\|v\|_1$, and for all $k \geq 4$, we have $\frac{dh}{dv_2} \geq \frac{dh}{dv_k}$.*
5. *For $v \in D$ with $v_1 \geq \frac{1}{2}\|v\|_1$, we have $h(v) \leq \|v\|_1 v_1 - \frac{3}{4}v_1^2$.*

Proof. Part 1 follows from the fact that h is a homogenous of degree two on $\ell^2(\mathbb{N})$. For part 2, we compute $\frac{dh}{dv_i}$. For $i = 1$, we see that $\frac{dh}{dv_1} = 2v_2$. For $i > 1$,

$$\frac{dh}{dv_i} = 2v_{i-1} - 2v_i + 2v_{i+1}.$$

Thus, for $k \geq 3$, we have

$$\frac{dh}{dv_1} - \frac{dh}{dv_k} = (2v_2 - 2v_{k-1}) + (2v_k - 2v_{k+1}).$$

This is positive, because v is a weakly decreasing sequence.

For part 3, define $w = (\|v\|_1 - v_2, v_2, 0, \dots) \in D$. Let

$$D_k = \{v \in D : v_{k+1} = 0\}.$$

[1]The condition that all v_i be non-negative is unnecessary. Because the terms are decreasing, if any term is strictly less than zero, the entire sequence would not be in $\ell^1(\mathbb{N})$. We include it as part of the condition for the sake of clarity.

Furthermore, define $e_i \in \ell^1(\mathbb{N})$ to be 1 in the ith position, and zero elsewhere. First suppose that $v \in D_k$ for some $k \geq 3$. From part 2, we know that

$$h(v + v_k(e_1 - e_k)) \geq h(v).$$

Moreover, $v + v_k(e_1 - e_k) \in D_{k-1}$. Iterating this process, we see that if $v \in D_k$ for some $k \geq 3$, then we have

$$h(v) \leq h(\|v\|_1 - v_2, v_2, 0, \dots) = 2\|v\|_1 v_2 - 3v_2^2.$$

Because $\bigcup_{k=2}^{\infty} D_k$ is dense in D, and both the left- and right-hand sides are continuous, we see that $v \leq 2\|v\|_1 v_2 - 3v_2^2$ for every $v \in D$.

For part 4, note that

$$\frac{dh}{dv_2} - \frac{dh}{dv_k} = 2[v_1 - (v_2 - v_3) - (v_{k-1} - v_k) - v_{k+1}].$$

If $k \geq 4$, the quantity subtracted from v_1 is less than $\|Lv\|_1 = \|v\|_1 - v_1$. Because $v_1 > \frac{1}{2}\|v\|_1$, the quantity $\frac{dh}{dv_2} - \frac{dh}{dv_k}$ must be non-negative.

Lastly, for part 5, we again consider $D_k = \{v \in D : v_{k+1} = 0\}$. For $v \in D_k$, note that part 4 implies $h(v + v_k(e_2 - e_k)) \geq h(v)$. Because $v_1 \geq \frac{1}{2}\|v\|_1$, we see $v + v_k(e_2 - e_k) \in D_{k-1}$. Iterating this process, we eventually reach

$$w := (v_1, \|v\|_1 - v_1 - v_3, v_3, 0, \dots).$$

Thus, if $v \in D_k$ for some $k \geq 4$, we have

$$h(v) \leq h(w) = -\|v\|_1^2 + 4\|v\|_1 v_1 - 3v_1^2 + 4\|v\|_1 v_3 - 6v_1 v_3 - 4v_3^2.$$

As both the left- and right-hand sides are continuous, and $\bigcup_{k=3}^{\infty} D_k$ is dense in D, this bound holds for all $v \in D$. From here it suffices to show that $\|v\|_1 v_1 - \frac{3}{4}v_1^2 - h(w) \geq 0$ for all v. To this end, note that

$$\|v\|_1 v_1 - \tfrac{3}{4}v_1^2 - h(w) = \tfrac{1}{4}\left(2\|v\|_1 - 3v_1 - 4v_3\right)^2 \geq 0.$$

Hence, $h(v) \leq h(w) \leq \|v\|_1 v_1 - \frac{3}{4}v_1^2$ as desired. □

Corollary 5.8. *The number of conjugacy classes in the group of upper-triangular matrices is bounded by*

$$\mathbf{k}(U_n) \leq p(n)\sqrt{n!}\; q^{\frac{n^2}{6} + \frac{n}{2}},$$

where $p(n)$ is the number of integer partitions of n.

Proof. Because

$$|\mathrm{Comm}(\mathcal{U}_n)| = q^{\binom{n}{2}}\mathbf{k}(U_n),$$

it suffices to show that

$$|\mathrm{Comm}(\mathcal{U}_n)| \leq p(n)\sqrt{n!}\; q^{\frac{2n^2}{3}}.$$

To this end, we stratify $\mathrm{Comm}(\mathcal{U}_n)$ by the Jordan type of the first matrix in each pair. From Theorem 5.6, we have

$$|\mathrm{Comm}(\mathcal{U}_n)| = \sum_{\lambda \vdash n} |\mathrm{Comm}(\lambda)| \leq \sum_{\lambda \vdash n} \sqrt{n!} q^{\frac{1}{2}(n^2 + h(\lambda'))}.$$

The sum is over $p(n)$ terms, so it suffices to show that each term is bounded $\sqrt{n!}q^{\frac{2}{3}n^2}$. In other words, it suffices to show that for every $\lambda \vdash n$, we have

$$h(\lambda') \leq \frac{n^2}{3}.$$

Let $D = \{v \in \ell^1(\mathbb{N}) : v_1 \geq v_2 \geq \cdots \geq 0\}$. Let $c \in \mathbb{R}$ such that $v_2 = c\|v\|_1$. Then from part 3 of Lemma 5.7, we have $h(v) \leq (2c - 3c^2)\|v\|_1^2$. This quantity is maximized at a value of $\frac{1}{3}\|v\|_1^2$, by taking $c = \frac{1}{3}$. For any $\lambda \vdash n$, we therefore have $h(\lambda') \leq \frac{n^2}{3}$, as desired. Thus, our upper bound for the number of pairs of commuting upper-triangular matrices is

$$|\mathrm{Comm}(\mathcal{U}_n)| \leq p(n)\sqrt{n!}\, q^{\frac{2n^2}{3}},$$

which implies the desired result. □

We are now ready to prove the following result, which trivially implies Theorem 1.2.

Theorem 5.9. *For every positive integer n, and every prime power q, we have*

$$\mathbf{k}(U_n) \leq p(n)^2 n!\; q^{cn^2+\frac{n}{2}}, \text{ where } c = \frac{20\sqrt{2}}{49} - \frac{41}{98} \approx 0.15886.$$

Proof. Because $|\mathrm{Comm}(\mathcal{U}_n)| = \sum_{\lambda,\mu\vdash n} |\mathrm{Comm}(\lambda,\mu)|$, it suffices to show that

$$|\mathrm{Comm}(\lambda,\mu)| \leq \sqrt{n!}q^{cn^2+\frac{n}{2}}.$$

Let δ and ε denote small positive quantities, each less than $\frac{1}{6}$, to be optimized later. We bound $\mathrm{Comm}(\lambda,\mu)$ in one of three ways, depending on the shapes of λ and μ. Let $v \in \ell^1(\mathbb{N})$ be defined by $v := \lambda'/n$, so that $\|v\|_1 = 1$.

First consider the case where $\left|v_2 - \frac{1}{3}\right| \geq \delta$. From Lemma 5.7, we use the bound $h(v) \leq 2v_2 - 3v_2^2$. As a function of v_2, the right-hand side attains its maximum at $v_2 = \frac{1}{3}$. Thus, for v_2 satisfying $\left|v_2 - \frac{1}{3}\right| \geq \delta$, the right-hand side is maximized at $v_2 = \frac{1}{3} \pm \delta$ with a value of $\frac{1}{3} - 3\delta^2$. Thus,

$$|\mathrm{Comm}(\lambda,\mu)| \leq \sqrt{n!}\, q^{\frac{n^2}{2}(1+h(v))} \leq \sqrt{n!}\, q^{\frac{n^2}{2}(\frac{4}{3}-3\delta^2)}.$$

Because $|\mathrm{Comm}(\lambda,\mu)| = |\mathrm{Comm}(\mu,\lambda)|$, we also have this bound if $w = \mu'/n$ satisfies $\left|w_2 - \frac{1}{3}\right| \geq \delta$. Henceforth we may assume that both v and w have their second entry in the open interval $(\frac{1}{3} - \delta, \frac{1}{3} + \delta)$.

Second, if $v_1 > \frac{2}{3} - \varepsilon$, then from Lemma 5.7, we use the bound $h(v) \leq v_1 - \frac{3}{4}v_1^2$. Because $\varepsilon \leq \frac{1}{6}$, we know that $v_1 > \frac{1}{2}$, and so v satisfies the hypotheses of part 5 of Lemma 5.7. The polynomial $x - \frac{3}{4}x^2$ is maximized at $\frac{2}{3}$, and so for $v_1 < \frac{2}{3} - \varepsilon$, we see that $h(v) \leq \frac{1}{3} - \frac{3}{4}\varepsilon^2$. Thus,

$$|\mathrm{Comm}(\lambda,\mu)| \leq f^\lambda q^{\frac{n^2}{2}(1+h(v))} \leq \sqrt{n!}\, q^{\frac{n^2}{2}(\frac{4}{3}-\frac{3}{4}\varepsilon^2)}.$$

By symmetry, this bound also applies if $\mu_1' \leq (\frac{2}{3} - \varepsilon)n$, so we may assume that both λ_1'/n and μ_1'/n are at least $\frac{2}{3} - \varepsilon$.

Lastly, consider the case where

$$\frac{\lambda_1'}{n} \geq \frac{2}{3} - \varepsilon, \text{ and } \left|\frac{\lambda_2'}{n} - \frac{1}{3}\right| \leq \delta,$$

and similarly for μ. We obtain an upper bound in the last case by disregarding commutativity, and using the bound

$$|\mathrm{Comm}(\lambda,\mu)| \leq F_\lambda F_\mu = f^\lambda f^\mu q^{n^2-n-\mathbf{n}(\lambda)-\mathbf{n}(\mu)}.$$

To maximize the exponent, we take $\lambda' = \mu' = \big((\frac{2}{3}-\varepsilon)n, (\frac{1}{3}-\delta)n, 1, 1, \ldots, 1\big)$. The exponent $n^2 - n - \mathbf{n}(\lambda) - \mathbf{n}(\mu)$ is therefore bounded above by

$$n^2\left[1 - (\tfrac{2}{3}-\varepsilon)^2 - (\tfrac{1}{3}-\delta)^2\right].$$

Thus, our three bounds on the exponent are given by

$$n^2\left[\tfrac{2}{3} - \tfrac{3}{2}\delta^2\right],\ n^2\left[\tfrac{2}{3} - \tfrac{3}{8}\varepsilon^2\right] \text{ and, } n^2\left[1 - (\tfrac{2}{3}-\varepsilon)^2 - (\tfrac{1}{3}-\delta)^2\right].$$

Taking $\varepsilon = 2\delta = \frac{4}{21}(5 - 3\sqrt{2}) \approx 0.1443$, it is easy to verify that all three of these quantities are equal to αn^2, where

$$\alpha := \frac{4}{49} + \frac{20\sqrt{2}}{49} \approx 0.65886.$$

Thus, we obtain the bound $\mathrm{Comm}(\mathcal{U}_n) \leq \sum_{\lambda,\mu} n! q^{\alpha n^2}$, proving that

$$\mathbf{k}(U_n) \leq p(n)^2\ n!\ q^{(\alpha - \frac{1}{2})n^2 + \frac{n}{2}} = p(n)^2\ n!\ q^{cn^2 + \frac{n}{2}}.$$

□

6. Lower central series

The aim of this section is to provide bounds on the number of conjugacy classes in groups in the lower central series of $U_n(q)$. The techniques used here do not overlap with the techniques of gap arrays introduced in the previous sections.

For a nonnegative integer $k \leq n-1$, let $U_{n,k}(q)$ denote the kth term in the lower central series of $U_n(q)$. Explicitly,

$$U_{n,k}(q) = \{A \in U_n(q) : A_{i,j} = 0 \text{ whenever } 0 < i - j \leq k\}.$$

In this section, we prove an upper bound on then number of conjugacy classes in $U_{n,k}(q)$. We first introduce some notation which we will use in the proof.

By $\mathcal{U}_{n,k}(q)$, we mean the $\mathbb{F}_q$-subalgebra of $\mathcal{U}_n(q)$ defined by

$$\mathcal{U}_{n,k}(q) := \{X : 1 + X \in U_{n,k}(q)\}.$$

Note that $\mathcal{U}_n(q) = \mathcal{U}_{n,0}(q)$. We omit the parameter q when the field is clear from context. For a finite group G, we define the *commuting probability* $\mathrm{cp}(G)$ to be the probability that two elements chosen uniformly at random from G commute. Clearly

$$\mathrm{cp}(G) = \frac{|\mathrm{Comm}(G)|}{|G|^2} = \frac{\mathbf{k}(G)}{|G|}.$$

For $A \in \mathcal{U}_a$ and $B \in \mathcal{U}_b$, let $T_{A,B} : \mathcal{M}_{a\times b} \to \mathcal{M}_{a\times b}$ by $T_{A,B}(X) := AX - XB$. Note the map $T_{\lambda,\mu}$ defined in Lemma 4.1 is shorthand for T_{J_λ, J_μ}. The following lemma is the key tool used in the proof of Theorem 1.3.

Lemma 6.1. *Let a and b be positive integers, and let k be an integer such that $0 \leq k < a + b$. Let $V = \{X \in \mathcal{M}_{a\times b} : X_{i,j} = 0$ whenever $j - i \geq a - k\}$. Then*

$$\mathrm{cp}(a+b, k) = \frac{1}{|\mathcal{U}_{a,k}|^2 \cdot |\mathcal{U}_{b,k}|^2} \sum q^{-\dim(\mathrm{Im}\, T_{A_1,B_1} + \mathrm{Im}\, T_{A_2,B_2})}$$

where the sum is over all $(A_1, A_2) \in \mathrm{Comm}(\mathcal{U}_{a,k})$ and all $(B_1, B_2) \in \mathrm{Comm}(\mathcal{U}_{b,k})$.

Proof. We begin with the vector space isomorphism $\mathcal{U}_{a+b,k} \to \mathcal{U}_{a,k} \oplus V \oplus \mathcal{U}_{b,k}$ given by

$$e_{i,j} \mapsto \begin{cases} (e_{i,j}, 0, 0) & \text{if } i,j \leq a \\ (0, e_{i,j-a}, 0) & \text{if } i \leq a, j > a \\ (0, 0, e_{i-a,j-a}) & \text{if } i,j > a. \end{cases}$$

Graphically, the isomorphism is shown in Figure 4.

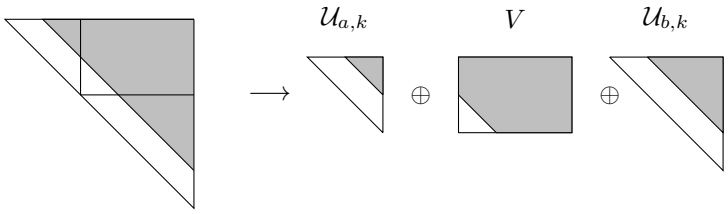

FIGURE 4. A graphical representation of the decomposition of $\mathcal{U}_{a+b,k}$ into $\mathcal{U}_{a,k} \oplus V \oplus \mathcal{U}_{b,k}$. The gray regions represent the cells which can be non-zero. The white regions represent the cells which must contain zeros.

The multiplicative structure on $\mathcal{U}_{a+b,k}$ can be pushed through this isomorphism to make $\mathcal{U}_{a,k} \oplus V \oplus \mathcal{U}_{b,k}$ an $\mathbb{F}_q$-algebra, and is given by

$$(A_1, X_1, B_1) \cdot (A_2, X_2, B_2) = (A_1 B_1,\ A_1 X_2 + X_1 B_2,\ A_2 B_2).$$

We count pairs of commuting elements in $\mathcal{U}_{a,k} \oplus V \oplus \mathcal{U}_{b,k}$. Note that (A_1, X_1, B_1) and (A_2, X_2, B_2) commute if and only all three of the following conditions hold:

(1) A_1 and A_2 commute in $\mathcal{U}_{a,k}$,
(2) B_1 and B_2 commute in $\mathcal{U}_{b,k}$,
(3) $T_{A_1,B_1}(X_2) = T_{A_2,B_2}(X_1)$.

For $i = 1, 2$, let $T_i = T_{A_i,B_i}$. We may count pairs of commuting elements in $\mathcal{U}_{a+b,k}$ by counting how many $(X_1, X_2) \in V \oplus V$ satisfy $T_1(X_2) = T_2(X_1)$, and summing over all $(A_1, A_2) \in \mathrm{Comm}(\mathcal{U}_{a,k})$ and all $(B_1, B_2) \in \mathrm{Comm}(\mathcal{U}_{b,k})$. Therefore,

$$\mathrm{cp}(a+b,k) = \frac{1}{|\mathcal{U}_{a+b,k}|^2} \sum \#\{(X_1, X_2) \in V \times V : T_1(X_2) = T_2(X_1)\},$$

where the sum is over

$$S := \{(A_1, A_2, B_1, B_2) : (A_1, A_2) \in \mathrm{Comm}(\mathcal{U}_{a,k}), (B_1, B_2) \in \mathrm{Comm}(\mathcal{U}_{b,k})\}.$$

Note that $\{(X_1, X_2) : T_1(X_2) = T_2(X_1)\}$ is the kernel of the map $\Phi : V \oplus V \to V$ defined by

$$\Phi(X, Y) := T_1(X) - T_2(Y).$$

Obviously, $\mathrm{Im}\,\Phi = \mathrm{Im}\,T_1 + \mathrm{Im}\,T_2$. By the rank-nullity theorem, we have

$$\mathrm{cp}(a+b,k) = \left(\frac{q^{\dim V}}{|\mathcal{U}_{a+b,k}|}\right)^2 \sum_S q^{-\dim(\mathrm{Im}\,T_1 + \mathrm{Im}\,T_2)}. \tag{6.1}$$

Recalling that $\dim \mathcal{U}_{a+b,k} = \dim \mathcal{U}_{a,k} + \dim V + \dim \mathcal{U}_{b,k}$, we see that

$$\frac{q^{\dim V}}{|\mathcal{U}_{a+b,k}|} = \frac{1}{|\mathcal{U}_{a,k}|\,|\mathcal{U}_{b,k}|}. \tag{6.2}$$

Putting together equations (6.1) and (6.2), we obtain the desired result. □

Lemma 6.2. *Let* $N_r(a,b) := \{(A,B) \in \mathcal{U}_a \times \mathcal{U}_b : \operatorname{rank} T_{A,B} \leq r\}$. *Then*

$$|N_{a,b}(r)| \leq p(a)p(b)\sqrt{a!b!}\; q^{\frac{(a-b)^2}{2}+r}.$$

Proof. Because $\dim\ker T_{A,B} = \dim\ker T_{C,D}$ whenever A is GL_a-conjugate to C and B is GL_b-conjugate to D, we may safely assume that A and B are in Jordan canonical form. From Lemma 4.1, we know $\dim\ker T_{\lambda,\mu} = \langle \lambda', \mu'\rangle$. Applying the Cauchy-Schwarz inequality followed by the classical arithmetic-geometric mean inequality, we obtain

$$\dim\ker T_{\lambda,\mu} = \langle \lambda', \mu'\rangle \leq \frac{\|\lambda'\|^2 + \|\mu'\|^2}{2} = \frac{a+b}{2} + \mathbf{n}(\lambda) + \mathbf{n}(\mu). \tag{6.3}$$

We stratify $N_{a,b}(r)$ by the Jordan forms of the pairs (A,B) which appear in $N_{a,b}(r)$. Let $S := \{(\lambda,\mu) : \langle \lambda', \mu'\rangle \geq ab - r\}$. The pairs of partitions in S are precisely the partitions indexing Jordan forms of pairs of matrices in $N_{a,b}(r)$. Thus,

$$|N_{a,b}(r)| = \sum_{(\lambda,\mu)\in S} F_\lambda(q)F_\mu(q).$$

Recall from (2.2) that $F_\lambda(q) \leq f^\lambda q^{\binom{a}{2}-\mathbf{n}(\lambda)}$, so

$$|N_{a,b}(r)| \leq \sum_{(\lambda,\mu)\in S} f^\lambda f^\mu q^{\frac{a^2+b^2}{2} - \left(\frac{a+b}{2}+\mathbf{n}(\lambda)+\mathbf{n}(\mu)\right)}.$$

Now applying (6.3) we get

$$|N_{a,b}(r)| \leq \sum_{(\lambda,\mu)\in S} f^\lambda f^\mu q^{\frac{a^2+b^2}{2}-(ab-r)} = q^{(a-b)^2/2+r} \sum_{(\lambda,\mu)\in S} f^\lambda f^\mu.$$

Noting that $f^\lambda \leq \sqrt{a!}$ and $|S| \leq p(a)p(b)$, the result follows. □

We now present the proof of Theorem 1.3. Define $\beta_0 := 0$, and for each positive integer m, define β_m inductively by

$$\beta_m := \tfrac{1}{4}\left(\beta_{m-1} - (1-2^{-m})^2\right).$$

A routine calculation shows that for all m, we have

$$\beta_m = -\frac{1}{3} - \frac{2}{3}\cdot 4^{-m} + 2^{-m} - 4^{-(m+1)}m.$$

Theorem 6.3. *Let* m *be a non-negative integer such that* $2^{-(1+m)} \leq \frac{k}{n} \leq 2^{-m}$. *Then*

$$\mathrm{cp}(n,k) \leq q^{\beta_m n^2(1+o_m(1))},$$

where $o_m(1)$ *is a function of* m *and* n *which, for each fixed* m, *tends to zero as* n *tends to infinity.*

Proof. We proceed with the proof that $\mathrm{cp}(n,k) \leq q^{\beta_m n^2(1+o_m(1))}$ by induction on m. When $m = 0$, we have $\frac{1}{2} \leq \frac{k}{n} \leq 1$, and so $U_{n,k}$ is abelian, proving that $\mathrm{cp}(n,k) = 1 = q^0$ as desired.

Now let $V = \{X \in \mathcal{M}_{a\times b} : X_{i,j} = 0 \text{ whenever } j - i \geq a - k\}$, and let

$$C^k_{a,b}(r) = \{(A_1,A_2,B_1,B_2) \in \mathrm{Comm}(\mathcal{U}_{a,k}) \times \mathrm{Comm}(\mathcal{U}_{b,k}) : \operatorname{rank}_V(T_{A_1,B_1}, T_{A_2,B_2}) = r\}.$$

From Lemma 6.1, we have

$$\operatorname{cp}(n,k) = \frac{1}{|\mathcal{U}_{a,k}|^2 |\mathcal{U}_{b,k}|^2} \sum_{j=0}^{ab-\binom{k+1}{2}} q^{-j} \left|C^k_{a,b}(j)\right|.$$

Pick some number r to be optimized later, and split the sum based on the relationship between r and j. First considering the case when $j \geq r$, we have

$$\sum_{j\geq r} q^{-j} \left|C^k_{a,b}(j)\right| \leq q^{-r} \sum_{j\geq r} \left|C^k_{a,b}(j)\right| \leq q^{-r} |\operatorname{Comm}(\mathcal{U}_{a,k})| |\operatorname{Comm}(\mathcal{U}_{b,k})|. \tag{6.4}$$

On the other hand, if $j < r$, we forget about the commutativity relation, and remember only that the rank of each map T_{A_i,B_i} must be bounded by r. Thus, for $j < r$ we have $C^j_{a,b}(k) \leq |N_{a,b}(j)|^2$. Combining this with (6.4) and Lemma 6.2, we obtain

$$\begin{aligned}
\operatorname{cp}(n,k) &\leq q^{-r} \operatorname{cp}(a,k) \operatorname{cp}(b,k) + \frac{1}{|\mathcal{U}_{a,k}|^2 |\mathcal{U}_{b,k}|^2} \sum_{j=0}^{r-1} q^{-r} |N_{a,b}(j)|^2 \\
&\leq q^{-r} \operatorname{cp}(a,k) \operatorname{cp}(b,k) + \frac{p(a)^2 p(b)^2 a!b!}{|\mathcal{U}_{a,k}|^2 |\mathcal{U}_{b,k}|^2} \sum_{j=0}^{r-1} q^{(a-b)^2+2j-r} \\
&\leq q^{-r} \operatorname{cp}(a,k) \operatorname{cp}(b,k) + \frac{p(a)^2 p(b)^2 a!b!}{|\mathcal{U}_{a,k}|^2 |\mathcal{U}_{b,k}|^2} q^{(a-b)^2+r-1}.
\end{aligned}$$

To optimize these parameters, we take $a = \lfloor n/2 \rfloor$ and $b = \lceil n/2 \rceil$, and

$$r = \left[\frac{1}{4}\left(\beta_{m-1} + \left(1-2^{-m}\right)^2\right) n^2\right], \tag{6.5}$$

where the square brackets denote the nearest integer function. In the calculations to follow, we omit the $[\cdot]$, as it complicates the computation, and does not contribute to the leading term in the exponent. The difference is absorbed into the $o_m(1)$ term.

Because $2^{-m} \leq \frac{2k}{n} \leq 2^{1-m}$, by inductive hypothesis, we have

$$\operatorname{cp}(a,k) \leq q^{\beta_{m-1} a^2 (1+o_m(1))} = q^{\frac{\beta_{m-1} n^2}{4}(1+o_m(1))}.$$

Note that $p(a)^2 p(b)^2 a!b! \leq q^{o(n^2)}$, and that

$$|\mathcal{U}_{a,k}| = q^{\frac{(a-k)^2}{2}(1+o_m(1))} = q^{\frac{1}{8}\left(1-\frac{2k}{n}\right)^2 n^2 (1+o_m(1))}.$$

Combining these facts, we obtain

$$\operatorname{cp}(n,k) \leq q^{-r+\frac{\beta_{m-1} n^2}{2}(1+o_m(1))} + q^{r-\frac{1}{2}(1-2^{-m})^2 n^2 (1+o_m(1))}$$

Substitiuting in (6.5), both exponents on the right-hand side become

$$\tfrac{1}{4}\left(\beta_{m-1} - (1-2^{-m})^2\right) n^2 (1+o_m(1)),$$

which is equal to $\beta_m n^2 (1+o_m(1))$. Thus,

$$\operatorname{cp}(n,k) \leq 2q^{\beta_m n^2 (1+o_m(1))} = q^{\beta_m n^2 (1+o_m(1))},$$

which completes the proof. □

Corollary 6.4 (Theorem 1.3)**.** *Let*

$$\gamma_m := \frac{1}{6} - \frac{13}{24} \cdot 4^{-m} + 2^{-(m+1)} - 4^{-(m+1)} m.$$

Then for every q,

$$\mathbf{k}(U_{n,k}(q)) \leq q^{\gamma_m n^2 (1+o_m(1))},$$

where $m = \left\lfloor \log_2 \left(\frac{n}{k}\right) \right\rfloor$, and $o_m(1)$ denotes a function which, for each fixed m, tends to zero as n tends to infinity.

Proof. Recalling that $\mathbf{k}(G) = \mathrm{cp}(G)\,|G|$, we see that Theroem 6.3 implies that

$$\mathbf{k}(U_{n,k}) \leq q^{\binom{n-k}{2} + \beta_m n^2 (1+o_m(1))}.$$

Combining this with the fact that $\binom{n-k}{2} = \frac{1}{2}\left(1 - \frac{k}{n}\right)^2 n^2 (1 + o_m(1))$, and also the fact that $\frac{k}{n} \geq 2^{-(1+m)}$, we obtain

$$\mathbf{k}(G) \leq q^{\left(\frac{1}{2}\left(1-2^{-(1+m)}\right)^2 + \beta_m\right) n^2 (1+o_m(1))} \leq q^{\gamma_m n^2 (1+o_m(1))},$$

where $\gamma_m = \frac{1}{6} - \frac{13}{24} \cdot 4^{-m} + 2^{-(m+1)} - 4^{-(m+1)} m$. □

7. Final Remarks

7.1. The original motivation for determining asymptotic bounds on $\mathbf{k}(U_n)$ was that such bounds could be used to give upper bounds on the number of p-groups of a given order. Specifically, Higman used his upper bound of $\mathbf{k}(U_n(p)) \leq p^{\frac{n^2}{4}}$ to show that there were no more than $p^{\frac{2n^3}{15}(1+o(1))}$ groups of order p^n [Hig]. In the same paper, he provided the lower bound of $p^{\frac{2}{27}n^3(1+o(1))}$ for the number of groups of order p^n, by explicitly exhibiting so many p-groups of class 2, but was unable to close the gap between the constants 2/27 and 2/15. He noted that a reduction in the upper bound for $\mathbf{k}(U_n(q))$ would result in a reduction for the upper bound on the number of p-groups, but that even proving $\mathbf{k}(U_n(q)) \leq q^{\frac{n^2}{12}(1+o(1))}$ could not completely close the gap with his methods.

In 1965, Charles Sims used a different argument to show that the number of groups of order p^n is $p^{\frac{2n^3}{27} + O(n^{8/3})}$, resolving Higman's original question [Sim] (see also, [BNV]).

7.2. Several other people have studied conjugacy in U_n since Higman first introduced the question. Gudivok et al. provide an algorithm to compute the number of conjugacy classes in U_n for arbitrary q, and computed the number for $n \leq 7$ [G$^+$].[2] Vera-López and Arregi constructed a different algorithm and were able to compute $\mathbf{k}(U_n)$ for arbitrary q and $n \leq 13$ [VLA03]. Moreover, they verified that for $n \leq 13$, the quantity $\mathbf{k}(U_n)$ was a polynomial in q with integer coefficients of degree $[(n^2 + 6n)/12]$, where $[\cdot]$ denotes the nearest integer function. In [PS], we verify that $\mathbf{k}(U_n)$ is a polynomial in q with integer coefficients of the same degree for $n \leq 16$. While there is evidence that $\mathbf{k}(U_n)$ may not be a polynomial in q for large n (see [HP, PS]), we believe that Higman's lower bound is asymptotically tight:

[2]The paper actually computed the number of conjugacy classes through $n = 9$, though the implementation had an error which yielded the wrong values for $n = 8$ and 9. The polynomials provided do evaluate to the correct values for $\mathbf{k}(U_n(2))$.

Conjecture 7.1. *For every prime power q,*

$$\mathbf{k}(U_n) = q^{\frac{n^2}{12}(1+o(1))}.$$

7.3. The proof of Corollary 5.8 shows that our bound is maximized for partitions with at most two columns. However, a straightforward inductive argument proves that if A has Jordan form J_λ, and λ has at most two columns, then there exists a gap array G such that $X_A C_{\mathcal{U}}(A) X_A^{-1} = C(G)$. Moreover, we can explicitly compute the size of this gap array, and determine that $\mathrm{Comm}(\lambda) = q^{o(n^2)}$.

7.4. Theorem 1.3 provides another proof that $\mathbf{k}(U_n) \leq q^{\frac{n^2}{6}(1+o(1))}$. Taking m to infinity, we see that γ_m approaches $\frac{1}{6}$. To make this precise, one must also show that the the term $o_m(1)$ does not grow too fast with m. In [Sof], we do this explicitly.

7.5. A *pattern group* is a subgroup of U_n defined by prescribing cells which are allowed to be non-zero (see [Isa]). The normal pattern groups in U_n were characterized by Weir, and have the property that if a cell is allowed to be non-zero, then so is any cell to the right of it or above it [Wei]. Due to this combinatorial structure we believe that similar techniques can be used to bound the number of conjugacy classes.

7.6. Regarding the bounds on $\mathbf{k}(U_{n,k})$, Lemma 6.2 has room for improvement. Specifically, when we wish to bound the number of matrices in $U_{n,k}$ which have Jordan type λ, we use F_λ, the number of such matrices in U_n. If we had estimates on the number of matrices in $U_{n,k}$ of a particular Jordan type, we could tighten the upper bound on $\mathbf{k}(U_{n,k})$. We hope to return to this problem in the future.

8. Acknowledgements

I would like to thank my Ph.D. advisor, Igor Pak, for introducing me to this problem, and for his help with early revisions of this paper. I would like to specially thank Persi Diaconis and Eric Marberg for pointing us to [Mar08], and Antonio Vera-López for directing us to [VLA92]. Their helpful comments and discussions were invaluable. I am also grateful to Adam Azzam and Alejandro Morales for their many helpful remarks, both mathematical and expository.

References

[BNV] S. R. Blackburn, P. M. Neumann, and G. Venkataraman, *Enumeration of finite groups*, Cambridge Univ. Press, Cambridge, UK, 2007.

[DI] P. Diaconis and I. M. Isaacs, Supercharacters and Superclasses for Algebra Groups, *Trans. AMS* **360** (2008), 2359–2392.

[G$^+$] P. Gudivok, Yu. Kapitonova, S. Polyak, V. Rud′ko, and A. Tsitkin, Classes of conjugate elements of a unitriangular group, *Kibernetika* **1** (1990), 40–48.

[HP] Z. Halasi and P. Pálfy, The number of conjugacy classes in pattern groups is not a polynomial function, *J. Group Theory* **14** (2011), 841–854.

[Hig] G. Higman, Enumerating p-groups. I. Inequalities, *Proc. LMS* **3** (1960), 24–30.

[Isa] I. Isaacs, Counting characters of upper triangular groups, *J. Algebra* **315** (2007), 698–719.

[Mac] I. G. Macdonald, *Symmetric Functions and Hall Polynomials*, (Second Ed.), Oxford Univ. Press, New York, USA, 1995.

[Mar08] E. Marberg, Constructing Modules of Algebra Groups, Undergraduate thesis, Stanford Univ. (2008).

[Mar11] E. Marberg, Combinatorial methods of character enumeration for the unitriangular group, *J. Algebra* **345** (2011), 295–323.

[PS] I. M. Pak and A. Soffer, Co-adjoint orbits and pattern groups, in preparation.

[Sag] B. E. Sagan, *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*, Springer, 2001.

[Sim] C. C. Sims, Enumerating p-groups, *Proc. LMS* **3** (1965), 151–166.

[Sof] A. Soffer, Enumerating conjugacy classes in unitriangular groups, and related combinatorial problems, UCLA Ph.D. thesis, in preparation.

[Sta] R. P. Stanley, *Enumerative Combinatorics*, Vol. 1 (Second Ed.), Cambridge Univ. Press, Cambridge, UK, 2011.

[VLA92] A. Vera-López and J. M. Arregi, Conjugacy classes in unitriangular matrices, *J. Algebra* **152** (1992), 1–19.

[VLA03] A. Vera-López and J. M. Arregi, Conjugacy classes in Sylow p-subgroups of $GL(n, q)$, *Linear Algebra Appl.* **370** (2003), 85–124.

[Wei] A. J. Weir, Sylow p-subgroups of the classical groups over finite fields with characteristic prime to p, *Proc. AMS* **6** (1955), 529–533.

[Yip] M. Yip, q-Rook placements and Jordan forms of upper-triangular nilpotent matrices, *DMTCS Proc.* (2013), 1017–1028.